Performance Localisation

Brendan Cody-Kenny*†, Michael O'Neill*, Stephen Barrett†

*Natural Computing Research & Applications Group
Michael Smurfit Graduate Business School
University College Dublin
Ireland
{brendan.cody-kenny, m.oneill}@ucd.ie

† Distributed Systems Group
School of Computer Science & Statistics
Trinity College Dublin
Ireland
stephen.barrett@scss.tcd.ie

arXiv:1603.01489v2 [cs.SE] 5 Sep 2016

***Abstract*—Performance becomes an issue particularly when execution cost hinders the functionality of a program. Typically a profiler can be used to find program code execution which represents a large portion of the overall execution cost of a program. Pinpointing where a performance issue exists provides a starting point for tracing cause back through a program.**

While profiling shows where a performance issue manifests, we use mutation analysis to show where a performance improvement is likely to exist. We find that mutation analysis can indicate locations within a program which are highly impactful to the overall execution cost of a program yet are executed relatively infrequently. By better locating potential performance improvements in programs we hope to make performance improvement more amenable to automation.

***Index Terms*—Performance; Mutation Analysis;**

I. Introduction

Software maintenance tasks, such as bug fixing and performance improvement are time consuming [1]. Once a bug is detected, it may be difficult to understand the existing code and design a fix. This is particularly difficult in larger more complex programs. To aid the diagnosis process, localisation techniques have been developed to highlight what code elements are particularly relevant to a functionality [2] or performance [3] defect. Finding where and how an issue manifests in source code can help indicate where a solution is likely to exist.

Improving program performance is frequently of secondary importance to improving functionality when developing software [4]. Focusing developer attention on functionality allows performance issues to manifest. Recent results indicate that improvement is mostly attempted when developers notice a clear improvement opportunity [1]. Performance improvement is also undertaken to alleviate performance bugs where program execution cost impinges on functional correctness [1]. Outside of these prominent scenarios, the implicit nature of how source code results in performance may further allow potential performance opportunities to go unnoticed. Where modern development practices recommend separation of concerns and reuse of code, it becomes increasingly unlikely that developers understand the performance characteristics of the API's and libraries their code depends on [5]. To aid performance issue detection, static analysis techniques have been developed for locating performance bugs [6] and bottlenecks [7] in code.

Profiling generally refers to measuring the execution cost of a program and is a frequently used technique for finding the location of performance "bottlenecks" in code. Profiling can be performed by instrumenting a program which increments a counter for each line of code each time it is executed. Additionally, a program can be executed with input of varying size to highlight what lines show an exponential increase in executing as input size increases [3]. Profiler techniques generally require developers must frequently trace back through a program to understand what code is contributing to a bottleneck [8]. Though code profiling can determine the location of a performance issue or bottleneck, it does not indicate what code change is required to improve performance. A performance improvement may not always be found at the same location. Finding a bottleneck does not always indicate the location of a solution to the bottleneck. This expresses the need for a technique which can determine more accurately where an improvement is likely to exist within a program.

As fault localisation [2] has been used to automate bug fixing [9], we seek similar methods for localising performance [3] to source code elements to benefit automated performance improvement [10], [11]. We seek localisation techniques which highlight code having the most effect on the overall execution cost of a program. Thus we are more interested in finding performance improvement opportunities than finding performance bottlenecks.

Our research question follows as:

- What performance localisation technique most accurately highlights locations of improvements?

Our hypothesis is that code locations which are particularly influential to program performance are likely good locations for finding performance improvements. To inspect this hypothesis, we consider to what extent code mutation can attribute performance to source code elements. Mutation has been used previously to find code locations which influence program performance [11] though this approach

only considers mutations which leave program functionality unaffected. In this paper we look for mutation locations which reduce execution cost regardless of effect on functionality. As the goal is to use mutation to find hints for where a performance improvement may exist program, we exclude any mutations which reduce execution cost while leaving functionality "correct".

We introduce 3 analysis techniques in section II based on two different types of mutation which highlight locations in code for their relevance to overall program performance. The first mutation approach deletes program statements (including any enclosed code) and measures the resulting savings in execution cost. The reduction in execution cost is attributed to deleted code. The second technique makes all possible changes to every modification location in a program and attributes performance change to these finer-grained modification points.

The intuition behind the use of mutation is that there are locations in a program which are "levers on performance" and have some disproportionately large control over execution cost. We are looking for code locations where small modifications produce a comparatively large change in overall program execution cost. These mutation-based approaches shift focus towards locations in code where modifications are likely to alleviate a bottleneck, performance bug or even some more modest improvement opportunity.

Motivating example. The canonical example involves a variable which is initialised early in program execution and determines how many times a loop executes later in the program. The execution cost of this variable initialisation is low as the line is only executed once. However, a large amount of the overall execution cost of the program can be attributed to this initialisation if the variable is later used as a condition for how many times a loop executes.

To further illustrate the point we use a simple BubbleSort algorithm with a redundant outer loop added as seen in Fig. 1a. The additional outer loop does not change the program semantics but causes the BubbleSort algorithm to be needlessly iterated over a second time. A profiler will give this outer loop a very low value in terms of execution cost as can be seen in Fig. 1b, therefore taking attention away from a prominent performance improvement opportunity.

The execution count in Fig. 1b shows the number of times each statement is executed as a percentage. When an array of size 10 with all elements in reverse order is passed as `a`, lines 4 & 5 are both executed 200 times. The execution count for each statement will vary depending on the distribution of values within the array. An array of 10 values with a different ordering will produce a different execution count profile and can change the ranking of each statement with respect to the others. If a fully sorted array is passed then statements 5, 6 and 7 will not be executed. A reverse sorted array executes each line the maximum number of times possible and is expected to give the same ranking of statements as input size is increased. If profiling was used in this case to guide automated performance improvement [10], it would appear to decrease the chances of finding this performance improvement as effort is spent modifying other locations.

In contrast, deletion analysis shows how much of the program execution cost is attributable to the outer loop. Fig. 1c shows the amount of execution cost that is saved when a statement (including any sub-statements) is removed. Execution cost savings are a percentage of the overall cost of executing the program. Note that as deletion analysis removes a statement inclusive of any sub-statements, percentages are cumulative. When statement 2 is removed the entire body of the method is removed, and so close to 100% of the execution cost is saved. Statement 2 receives a marginally larger percentage and is ranked ahead of statement 3. Deleting line 6 will result in a program which does not compile, in this case the line receives the execution cost saving from its parent statement as all nodes within a statement (and any sub-statements) are given the same value initially.

Contributions. In this paper we inspect the use of mutation to indicate the location of performance improvements in programs. We inspect two types of mutation:

- Deletion analysis takes advantage of the hierarchical nature of source code. Source code statements (and any child statements) are deleted and the resulting program variant is executed.
- Exhaustive mutation analysis makes all possible changes to each modification point in a program. For each code location, we can generate a set of program variants, one for each possible change at that location. In other words, we generate all possible first order mutants [12] for each modification point in a program.

Using the results of these mutation techniques the characteristics of program variants are analysed for each location in the program. Using simple heuristics across the summarised information for each location, it is possible to determine the likelihood of a location being important for improving performance. We evaluate profiling against 3 different analysis approaches in section II for localising performance based on the previously mentioned mutation techniques:

- Under deletion analysis (subsection II-B), the difference in performance between the original and variant programs is attributed to all code which was deleted. Every statement in a program is tested in this way, giving code inside inner loops a lower value than their enclosing looping constructs .
- Under exhaustive mutation (subsection II-C), the number of times a program variant shows reduced execution cost is divided by the number of times a mutation results in a compiled program.
- Occasionally there is no possible single point mutation which will produce a compilable program. In this scenario we use the results of deletion analysis to fill

```
void sort(Integer[] a, int length){                    1
 for (int h = 0; h < 2; h++) {                         2
  for (int i=0; i < length; i++){                      3
   for (int j=0; j < length - 1; j++){                 4
    if (a[j] > a[j + 1]){                              5
     int k=a[j];                                       6
     a[j]=a[j + 1];                                    7
     a[j + 1]=k;                                       8
}}}}}                                                  9
```

(a) "BubbleLoops" problem: Bubblesort with an extra redundant outer loop

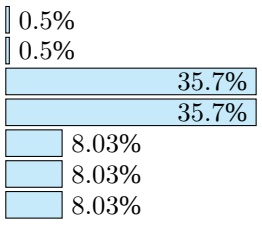

(b) Profiler: Execution frequency for each statement

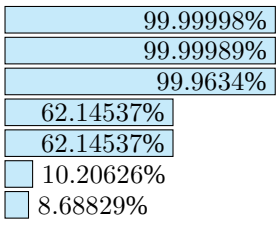

(c) Deletion Analysis: Execution savings when a statement (including sub-statements) is deleted

Fig. 1: BubbleLoops problem and profiles

in the gaps where exhaustive analysis was not able to glean any information (subsection II-D).

We evaluate these approaches on a set of test problems and find:

- profiling achieves the highest accuracy of all approaches on specific nodes, but does not generalise across our problem set (Table III)
- mutation analysis can, on average, better highlight locations of possible performance improvements in code (Figure 3)
- that there is a trade-off between the amount of computation required for each approach and the accuracy (subsection IV-A)

II. Performance Localisation Techniques Studied

In this section we describe the localisation techniques that are compared in our evaluation. Though many static analysis techniques exist for detecting performance issues, we compare with a profiling approach as recently used for automatic performance improvement [10].

A. Profiling

Our approach to profiling is relatively fine-grained to other approaches which may measure, for example, elapsed time for method execution. We measure the number of times each statement in a program is executed. For each source code statement, as defined in the Java language specification [13], we add an instrumentation statement as demonstrated in Figure 2. Each instrumentation statement consists of a function call with a program identifier and the line number for that location. When the instrumented program is run an execution count for each line is gathered.

B. Deletion Analysis

Deletion analysis was designed in an attempt to shift focus from bottlenecks towards code which has some influence over performance. A program has a statement removed and the resulting program variant is evaluated. When a statement contains sub-statements, for example a "FOR", "WHILE" or "IF" statement, the inner block statement and all sub-statements are removed also. The hierarchical structure of imperative code is made accessible by using Abstract Syntax Tree (AST) parsers [14] and is also used in scoping variables in code. Statements are removed in order of their appearance

```
void sort_5(Integer[] a, int length){
 ASTInstrumenter.recordExecution(5,0)
 for (int h = 0; h < 2; h++) {
  ASTInstrumenter.recordExecution(5,2)
  for (int i=0; i < length; i++){
   ASTInstrumenter.recordExecution(5,3)
   for (int j=0; j < length - 1; j++){
    ASTInstrumenter.recordExecution(5,4)
    if (a[j] > a[j + 1]){
     ASTInstrumenter.recordExecution(5,5)
     int k=a[j];
     ASTInstrumenter.recordExecution(5,6)
     a[j]=a[j + 1];
     ASTInstrumenter.recordExecution(5,7)
     a[j + 1]=k;
    }
   }
  }
 }
 ASTInstrumenter.recordExecution(5,1)
}
```

Fig. 2: Instrumented "BubbleLoops" problem, counting lines 0 to 7 for program variant ID 5

in a breadth-first approach as per AST structure. Outer loops are removed before inner loops, with the most nested code being removed last.

Deletion analysis exploits the ordered and hierarchical structure of imperative code as execution cost is attributed to statements which appear earlier in the code and to statements higher in the hierarchy. For example, the body of a "FOR" loop which may contain many statements can be considered the child of a "FOR" statement. The parent "FOR" statement is attributed the execution cost of all child statements, however deeply nested, by deleting the "FOR" statement including all child statements and measuring the execution cost reduction of the variant program when compared with the original.

Deletion analysis may not always be applicable for every statement in a program. Consider statement 6 in Fig. 1a which initialises the variable `k`. Deleting this line will result in a program which does not compile and which we cannot evaluate. Our choice in this scenario is to either attribute zero to this statement or attribute a ranking by way of its enclosing statement higher in the hierarchy (in this case it would receive a value of 62.15 from statement 5).

C. Exhaustive Mutation Analysis

Deletion analysis is not always able to directly attribute information about the cost associated with all statements in a program. The approach also gives all code elements within a statement the same value. We are thus motivated to inspect mutation of code elements at a finer level than statement. This approach will give different rankings to individual code elements within a statement and for statements which could not be evaluated under deletion. It is specifically designed to perform a type of performance sensitivity analysis by way of mutation. For example, any valid change to variable initialisation or the loop condition in the outermost loop (statement 2 in Fig. 1a) is likely to show a pronounced change in the execution cost of the resulting program.

A set of program variants is produced by repeatedly replacing a node with every other valid alternative node as found in the original program. A program P, is made up of a set s of all elements in the program. Let s' be a set of clones of all elements in s. For each element l in s in the program P, a variant program can be generated by exchanging l for each of the elements in s'. In other words, exhaustive mutation analysis takes all code elements in a program and replaces them with all alternatives. Alternative code elements are gleaned from the program itself but we add all language-defined [13] operators regardless of whether they are contained in the program. While this appears to produce a large number of costly program evaluations, in practice not all variant programs are compilable or evaluatable as shall be seen in subsection IV-A.

Each node in a program is attributed a value by taking the number of times a modification resulted in a program with reduced execution cost divided by the number of times a modification resulted in a compilable variant as written in Equation 1.

$$NodeVal = \frac{N_{executionreduction}}{N_{compiled}} \quad (1)$$

Equation 1: Exhaustive Analysis gives each node a quotient value of the number of times the execution cost is reduced, divided by the number of times a compilable program is created.

Example results of values attributed by exhaustive mutation analysis to statement 2 in Fig. 1a `for (int h = 0; h < 2; h++) {` are shown in Table I. It is not possible to modify some nodes in the AST as listed in the table by "-". Where no mutation can be produce a compilable program variant, the value is 0.

TABLE I: Exhaustive mutation analysis example on a single line of code taken from the BubbleLoops problem

Node Num	Textual representation	Value
1	`for (int h...`	1.0
2	`int h=0`	-
3	`h=0`	-
4	`h`	.6
5	`0`	.7
6	`h < 2`	-
7	`h`	.16
8	`2`	.85
9	`h++`	-
10	`h`	0

D. Exhaustive and Deletion Combined

Even though exhaustive mutation analysis makes program modifications at a sub-statement level, it is still possible that no single element modification is able to produce a compilable, and hence evaluatable, program variant. Where no compilable program can be produced by modifying a particular location in a program we are missing information about it's relevance. This can occur where variable scoping prevents replacement by any other variable. To alleviate this issue, we use the results of deletion analysis to "fill the gaps" in the results of exhaustive mutation analysis where no single change produced a compilable program.

III. Methodology

To evaluate the analysis techniques, we use a set of problems with known performance improvements. By comparing programs with their variants which contain known performance improvements we can find what code elements differ. The nodes that differ between programs are considered improvement opportunities in the inefficient versions of

a program. These “improvement” nodes are of highest importance and receive the highest rank among all nodes in the program. The highest ranked and therefore most important nodes are those which are required to change to improve the performance of a program. We apply performance localisation techniques to these problems and gather node rankings. We compare the node rankings to our idealised ranking to determine which mutation technique is most accurate.

A. Program Evaluation Measures

We measure performance in terms of execution cost. For profiling, this is the number of statements executed when the program is run. For mutation approaches, we use a more fine-grained measure of execution cost by counting the number of byte-codes executed by the JVM [15].

B. Problem Set

We use a variety of sort implementations and a Huffman code-book (or “dictionary”) generation implementation to observe how profiling and mutation analysis can find improvement locations. A long-form code listing of all programs in our problem set is available [1].

Table II lists the implementations used and provides descriptive measures for each program as well as improvement types:

- **LOC** refers to the number of lines of code in the program
- **AST nodes** refers to the number of modification points in each program when it is parsed into an Abstract Syntax Tree representation [14].
- **Imp Nodes** refers to the number of nodes or locations in the program which need to be changed to achieve an improved version of the program. Although there are a number of different code modifications which can yield an improved variant of a program, we use the improvements which give the greatest reduction in execution cost with the smallest number of modifications. Modifications which reduce program functionality the least when applied individual are also favoured. We thus use the improvements that are the ’easiest’ or most probable set of modifications to be found with a search algorithm. This includes nodes involved in multiple known improvements.
- **Improvement** refers to the largest percentage improvement in execution cost known for each program [16].
- **Improvement Types** gives a high level description of the known improvement types for each program. For loop unrolling, the important node is the containing block statement.

[1] https://www.scss.tcd.ie/~codykenb/locoGP-ImprovementsFound.html

C. Test Cases

The execution cost of the test problem implementations we consider are affected by input size and distribution. We use a range of input sizes and distributions to ensure the profile is general. Input array size for sort implementations ranges from one to ten values. The distribution includes random, fully and reverse sorted ordering. For the Huffman code-book problem five different test cases which include arrays with repeated sequences and those without any repeated character.

D. Comparing Localisation Techniques

We have an idealised “best” ranking of nodes which put nodes involved in some improvement at the top which we term “improvement” nodes. These top ranked improvement nodes are required to change to produce a known improvement in each program. Fractional ranking is used as all nodes in a statement jointly share a given ranking. E.g. if two nodes share first place, then both nodes are given the ranking of "1.5" as this will be the ranking of the nodes on average if they are selected randomly.

For each program, each localisation technique produces a ranking for all nodes. Only a small proportion of these nodes are required to change to improve the performance of a program. The closer an improvement node is to where it should be in the idealised ranking is used as our measure of “accuracy”. The distance an improvement node is from where it should be in our idealised ranking is what we use as our “ranking error” measure. We normalise the ranking error for each node by dividing it by the number of nodes in the overall program to find at what percentile the node is placed.

For each technique we compare the percentile ranking error of each important node across all problems. This gives us 48 important nodes across all problems for comparison. We also do pair-wise comparison between techniques to be sure there is a statistically significant difference between them directly. We find the different between the approaches by subtracting the percentile rankings of the important nodes. We use a bootstrapping technique to analyse these differences. We sample randomly from these differences 100 times, with replacement, and calculate the average. This is repeated 100 times. This bootstrap approach gives an estimator for mean and approximate 95% confidence interval are given by the 0.015 and 0.975 quantiles.

We further summarise results be looking at nodes ranked in the upper 50 percentile of all the nodes represent instances where profiling has accurately highlighted the location of an improvement. We chose the 50 percentile as a simple way to show how the nodes in a program can be segregated. How the ranking of nodes is used to infer importance is dependent on the probability distribution generated from this ranking. The 50 percentile represents the median of node rankings with nodes in the upper half being considered

TABLE II: Problem Improvement Overview

Problem Name	LOC	AST Nodes	Imp Nodes	Imp	Improvement Type
Insertion Sort	13	60	3	9%	Loop unrolling
Bubblesort	13	62	5	45%	Redundant Traversal (exclude sorted portion)
BubbleLoops	14	72	8	71%	Redundant Traversal (exclude sorted portion)
Selection Sort 2	16	72	1	11%	Removed redundant increments during tests
Selection Sort	18	73	1	2%	Removed redundant array access
Shell Sort	23	85	3	5%	Various changes in increment size
Radix Sort	23	100	3	3%	Reduced iteration, comparison with 0
Quick Sort	31	116	2	54%	Reduced iterations, remove tests
Cocktail Sort	30	126	2	15%	Cloned and perforated loops (loop unrolling) Redundant Traversal (exclude sorted portion)
Merge Sort	51	216	1	5%	Remove redundant array clone
Heap Sort	62	246	2	41%	Remove redundant array access and assignment
Huffman Code-book	115	411	5	43%	Same as Bubblesort

$$PercentRankError = \frac{R_i - R_l}{N_{total}} \quad (2)$$

Equation 2: Percentile Ranking Error measure calculated for each improvement node.

more important than those in the lower half. Nodes which have a ranking in the upper 50 percentile of all nodes represent instances where profiling can be said to have been "accurate". As increasing the ranking of one node reduces the ranking of another, where an improvement node is in the lower 50 percentile of all nodes then the technique can be said to be "deceived".

A normalised percentile ranking error is the distance a node is ranked from its ideal ranking, divided by the number of nodes in that program (Equation 2).

IV. Results

The four performance localisation techniques, Profiling, Deletion, Exhaustive and Deletion with Exhaustive gap filling (Ex & De), are compared for accuracy in Table III.

We show a split at the 50 percentile to make the point that using a probability distribution over these accuracy values will result in some cutoff point where nodes below will receive lower importance and those above will receive higher importance (in comparison to a scenario where all nodes have the same ranking or importance). We can conceive of importance being only those nodes which are in the top 1% of all nodes. In such a scenario, profiling is the only approach which would designate any node as important. Profiling would be considered best in this scenario but would only highlight a single improvement node as important. The lower we place the threshold for importance as a percentile, the larger the combinations of those nodes become. The more of the important nodes we want to include as important, the more program nodes we must consider. To include all important nodes we must consider all nodes in the program, which does not help us reduce the number of nodes worth considering important. The more nodes we consider, the exponentially more combinations we need to consider.

When interpreting Table III we consider Exhaustive with Deletion (Ex & De) to be the best as this approach places the largest number of improvement nodes in the upper 50 percentile. The three mutation-based approaches also put a majority of the improvement nodes in the upper half of all nodes.

TABLE III: The accuracy of performance localisation techniques.

Accuracy	Profiler	Deletion	Exhaustive	Ex & De
99-100%	1	0	0	0
90-99%	7	8	9	11
80-90%	7	2	9	6
70-80%	3	10	7	5
60-70%	3	6	5	9
50-60%	2	4	2	5
40-50%	2	3	4	1
30-40%	6	2	5	3
20-30%	10	5	1	6
10-20%	2	2	0	2
0-10%	5	6	6	0

We further show a pair-wise comparison of the approaches using a bootstrap statistical technique (as described in subsection III-D) over the differences of percentiles for each improvement node. Figure 3 shows the difference between Profiling and Deletion, Deletion and finally Exhaustive and Exhaustive with Deletion. On average, improvement nodes are ranked roughly 2.75 percentage points higher under Deletion analsis when compared with a Profiler, 6.25 percentage points higher under Exhaustive analysis when compared with Deletion, and 3.6 percentage points higher still when using Exhaustive with Deletion.

Figure 3 also cross validates our evaluation as the differences in improvement node percentiles correlate with the ordering (though not magnitude) of which techniques are more accurate than others in Table III. The difference between the number of improvement nodes ranked in the upper half of all nodes as shown in Table III (Deletion ranks more nodes in upper half than Profiling, Exhaustive more than Deletion, and Exhaustive & deletion gap filling more than Exhaustive alone).

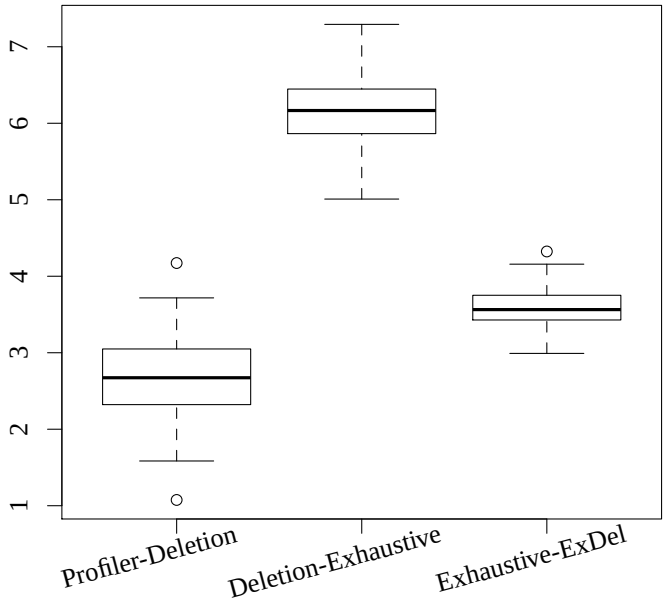

Fig. 3: Comparison of the **differences** between node percentile rankings for the four different approaches

Table IV shows descriptive summary values for each technique (Long form results are available [17].

Nodes most accurate shows for how many nodes each technique is the most accurate of all techniques. Profiling has the highest accuracy values on the most nodes (13 out of a total of 48 important nodes).

Nodes least accurate sums the number of improvement nodes each technique attributes the lowest ranking of all techniques. Deletion analysis gives the lowest ranking to the most nodes when compared with all other techniques.

Nodes ranked in upper half and **Nodes ranked in lower half** show a sum of the number of nodes ranked above and below the 50 percentile respectively.

Problems with only accurate nodes counts the number of problems which do not contain any improvement nodes ranked in the lower 50 percentile.

Problems majority nodes deceived shows a less stringent count of the number of problems where a majority of the improvement nodes are ranked in the lower half of all nodes. Where a majority of nodes are given a low ranking a technique can be said to be "deceived" as to the location of an improvement.

Best on Problems refers to when a technique has gives a majority of nodes the highest ranking.

Although profiling is accurate on 23 nodes across 8 problems, it is also deceived on a majority of the improvement nodes for 6 of the problems. It did however perform better than any other approach on 3 of the 12 problems including the Huffman Code-book problem which is the largest in our test set.

The use of Exhaustive analysis with deletion refinement ("Ex & Del" in table IV) was least deceived of all techniques across all nodes, with only 12 nodes lower than the 50 percentile of all nodes. It was deceived on at least 1 node in 7 of the 12 problems and was deceived on the majority of important nodes in 2 of the problems. It also has the highest accuracy on 12 of the 48 important nodes. It performs the best across 4 of the problems.

In these results we assume that if an approach is deceived on a majority of important nodes in a problem, it is likely that it will take longer to find an improvement as GP modifies other locations in the program. If half or more of the nodes are ranked highly, then it is likely that the approach will help GP find at least one of the possible improvements more quickly.

We consider a technique's ability to avoid being deceived as being more important than being the most accurate. We expect that there is some threshold value below which the use of a technique to guide a search process would lower the chances of finding an improvement. This threshold value is likely to have a polarising effect on a search algorithm. Effort spent modifying irrelevant nodes is effort that is not spent on important nodes. Due to this, we can hypothesise that a search algorithm would be delayed in finding performance improvements when focusing too much search effort on irrelevant nodes.

This is most obviously exemplified in the hand-crafted "BubbleLoops" problem, where an extra redundant outer loop has been added to Bubblesort. Profiling attributes a very low ranking to locations where simple changes which would half the execution cost of the program. Other examples which were not specifically crafted to be deceptive problems include Selection 2, Selection, Shell, Radix and Cocktail sort.

A. Computational Cost of Analysis

Profiling is the cheapest analysis to perform as instrumentation only need be performed once and a single execution of a program is needed to gather results. Even if we use profiling to find each statement's sensitivity to input size [3], we may only need use a small number of test cases to find this information. As evaluation time dominates, we use this as our measure of computational cost. We take profiling to cost a single evaluation.

TABLE IV: Summary

	Profiler	Deletion	Exhaustive	Ex & Del
Nodes most accurate (out of 48)	**13**	10	**13**	12
Nodes least accurate	15	19	8	**6**
Nodes ranked in upper half	23	30	32	**36**
Nodes ranked in lower half	25	18	16	**12**
Problems with only accurate nodes (out of 12)	4	4	**5**	**5**
Problems majority nodes deceived	6	**2**	**2**	**2**
Best on Problems	3	1	**4**	**4**

When a program is mutated, the possible evaluation categories a variant program may fall into:

1) Not Compilable
2) Infinite Loop
3) Run-time Error where functionality & execution cost differs from original program
4) Functionally degraded where functionality & execution cost differ
5) More expensive (execution cost only differs from original program)
6) Identical to original program in terms of functionality and execution cost measures
7) Less expensive in terms of execution cost

Previous results indicate that a large proportion (71 - 84%) of variant Java programs do not compile [16] although these values are found under a wider range of mutations than considered here (statement cloning is allowed).

All statements in a program can be legally deleted per the Java syntax due to the context insensitive nature of it's grammar regarding statements. Deletion analysis, as we implemented it, requires instrumentation for every variant program. As we create a variant program by deleting each statement the number of evaluations required is almost linear to the number of statements in a program. In practice it is slightly less than linear as deleting some lines of code results in a program variant which does not compile and does not need to be evaluated. Evaluating whether a program does not compile is quicker relative to the time it takes to fully evaluate a runnable program. Deletion based localisation strikes a balance between being relatively accurate across many problems and having an execution cost linear with program size.

On the face of it, exhaustive analysis for a program containing n elements gives n
! combinations. Evaluation is not required where a single point mutation is not possible due to the Java type system. The existince of duplicate code elements in a program reduces the number of variants that need to be evaluated also. Less still, are the number of programs which compile. Unfortunately we can not exclude programs with infinite loops[2] and runtime errors [3]. In any case, exhaustively mutating all code elements with all other elements in a program is practical for the relatively small programs in our test set. Although many replacements are not possible due to language typing constraints as enforced by the AST representation used, exhaustive mutation remains expensive, requiring the attempted replacement of every node with every other node. Many replacements will result in programs which can be quickly found to not compile, and therefore do not incur the comparatively large evaluation cost of repeat variant program execution with several different test input values.

B. Threats to validity

The main threat to validity of our results is the size of the problem set with the concern being that our results do not generalise outside this set. This issue is of particular concern due to the limited variety of program type in our set; all but one of our test programs implements a sorting algorithm. Though the problem set of Sort and Huffman Codebook problems appears to be varied enough to make ranking improvement nodes highly across all problems currently unattainable, there remains a potential issue that the approach of exhaustive mutation and deletion analysis has been specialised to the algorithms in our problem set. Adding problems to the test set with particular attention paid to choosing a wider variety of problem types would reduce this concern. The length of programs is relatively small which calls into question how accuracy is affected when analysis is performed over much larger programs. As we use a sum total of execution cost it may be more difficult to measure how a mutation affects the overall cost.

The important nodes listed in our tables are sometimes part of multiple possible improvements. There are dependencies amongst some of the nodes where modifications must be made in a certain sequence to yield an improved program

[2]As we cannot determine for how long a program will execute, we somewhat arbitrarily choose a practical timeout of 2.5 times the program's execution.

[3]We could conceivably estimate the characteristics of a program which shows a run-time error on certain input values but runs successfully on other input values, for example, smaller input values or values which are already sorted which cause less of the code to be executed are less likely to cause run-time issues. This information is telling in itself, and could also represent an additional dimension to location analysis.

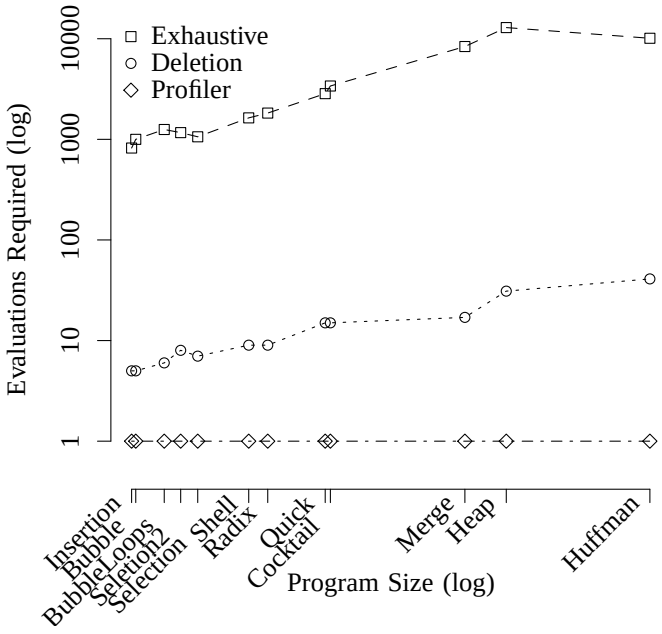

Fig. 4: Comparison of the cost of analysis between Exhaustive, Deletion and Profiling. Profiling is flat, requiring only one evaluation. Deletion is linear with program size. Exhaustive is exponential in relation to the number of AST nodes free to be modified.

making some improvements easier to find than others. Not all nodes are equal, given that a change in some may produce low functionality programs and are dependent on other modifications. As not all nodes are equal in terms of dependencies a simple summation summary may not appropriately capture a localisation techniques accuracy. If a majority of important nodes in a program are highly accurately identified it may not improve search where these nodes depend on one specific node which has unfortunately been misidentified. The "importance" of nodes is thus not uniform. This concern can be addressed by a closer inspection of how “difficult” each improvement is to achieve. If an improvement requires multiple changes to the program it can be said to be more difficult to find than an improvement requiring only a single modification.

V. Discussion

The major advantage of Profiling is the relatively low computational cost required. A single run of an instrumented program is enough to profile. Deletion analysis requires a program execution for each statement in a program though is less deceived on average than a profiler. Exhaustive mutation is more accurate still but also considerably more expensive to perform. The major disadvantage of using mutation is the high computational cost to perform localisation.

The problem size we use is relatively small and there is a potential limitation especially with exhaustive mutation regarding scalability. A potential solution might be to use a hybrid of approaches. Deletion analysis could be used initially to find what statements influence execution cost the most. We could only perform deletion analysis for code blocks which appear to be worth it. If removing the outermost loop reduces execution cost by some small fraction of overall cost, it may not be worth deleting and executing further nested statements. At some depth in the program subtree deletion analysis can be skipped where execution cost savings are negligible. Once deletion analysis has identified the most influential lines of code, exhaustive mutation can then be used sparingly to only distinguish between nodes within highly influential statements. Such an approach would further exploit the hierarchical structure of source code.

We use our results to say that the location of a performance bottleneck, as typically found using a profiler, does not always highlight potential performance improvements. When a performance improvement receives a low ranking, a search algorithm such as genetic programming will be less likely to find the improvement than had there been no node ranking at all.

The cost of performing mutation can be offset in scenarios when mutation is performed for other purposes such as mutation testing [18] or genetic improvement [9], [19]. Our results in this paper show that it is worth further attempting to further exploit the information generated by repeatedly executing mutated programs. Our main use case for this approach is as a guide for Genetic Programming (GP) to find performance improvements [16]. As mutation is the main driving force of the GP search process, performance localisation can potentially guide GP to performance improvements more quickly.

VI. Related Work

If we consider multiple versions of large programs as program variants then we can say that many modified versions of a program have been used to find performance issues on large scale software [20]. Using mutation to create many program variants has long been studied for software testing [18] and has recently been inspected for understanding the robustness of software [21], [12].

Closer to our approach is the use of mutation to discover “deep parameters” or locations where a modification in code relevant to program performance [11]. A deep parameter is a program mutation which affects program performance but not functionality. If program functionality changes in any measurable way, per an available test suite, the code location which was modified is removed from consideration as an interesting location for performance improvement. In contrast, our work shows that there is value in considering the location of mutations which degrade functionality but crucially also reduce execution cost.

The implicit nature of performance means it can be difficult to understand the nuances and interactions between program source code, input values and execution environment. This is especially true in many JavaScript environments where execution of the same code can differ widely [5]. A code change can provide better performance in some environments but can reduce performance in others. A reasonable amount of improvements actually reduce performance [5]. This may be due to a perceived opportunity that is not actually validated with any empirical evidence or an improvement that is beneficial in one environment but increases execution cost in another. This makes the case for using more automated tooling to localise performance and aid program understanding.

Many traits of a program (or "program spectra") can be measured to decompose a program internally [22]. A prominant approach to aid program understanding is to count the execution frequency of source code objects which is commonly referred to as profiling [23]. Profiling is widely used to find performance "bottlenecks" in code and highlights where the symptoms of performance issues can be observed [7], [8], [24], [25]. Many works build on the basic concept of profiling by using a range of input values to decompose and seperate the parts of a program and has been shown to be scalable to large numbers of inter-dependent input values [26]. Input sensitive profiling uses progressively larger sized input values to highlight what lines of code have a particularly acute response to increased program input size [3]. A similar approach has been used with success to guide GP [10] showing the importance of performance localisation.

Randomised input values can produce different executation traces which are used to isolate the root-cause of a performance issue. By finding a baseline expected performance on some input, subsequent input values can expose instances where performance is outside baseline. Root-cause can be determined by correlating system state against these anomalous instances [27]. Root-cause is pinpoonted by finding the "divergent point" where an anomalous execution deviates from expected. From this point to the point where a performance bug manifests can be considered suspicious. This approach was tested on large scale distributed systems.

Although there is no clear definition of what constitutes a performance bug we can generally say that a bug of this type is encountered when a program takes prohibitively longer than is expected or necessary to produce an output [1]. Some definitions use the existence of quadratic or cubic asymptotic execution response as input size increases as a clear indicator of a performance bug. Thus we feel that the definition of what constitutes a performance bug to be somewhat subjective. As we would like to have all programs execute as quickly as possible and the absence of a performance bug should not prevent us from attempting to improve program performance. A useful trait of profiling for improvement not specific to performance bug finding is that it provides a ranking of all code elements in a program as they contribute to program execution cost. A full ranking of code gives an ordering to what statements contribute the most without imposing a strict cutoff for what constitutes a "performance bug". Conversely, where a quadratic or cubic execcution response may be optimal for a specific algorithm and input, it would be unfair to classify problems which there are no known sub-exponential solutions [28] as performance bugs due to programmer error [1].

Static analysis is a lightweight alternative to dynamic analysis for finding performance issues. Static analysis appears to be more specific to certain types of performance issues[29], [30]. One advantage of identifying specific performance issues is that automatically fixing these performance bugs may be achieved by applying code changes which are known to frequently provide a fix [31]. In the current form of this work, when a performance bug is detected the unit of code marked as relevant to the bug is a (comparitively coarse-grained) function (or method in Java) [31].

Coarse-grained approaches which use a method or function as the smallest unit of code considered, appear more scalable for larger programs [6]. We see such approaches as complimentary where progressively less scalable approaches (such as exhaustive mutation) is only used after more scalable approaches have been used to broadly indicate what methods or libraries are associated with a performance issue.

VII. Conclusion & Future Work

We have shown how profiling is suited to finding performance bottlenecks and how this differs from locating performance improvements. Mutation analysis can be used to locate performance improvements and we show that even mutations which degrade functionality are worth analysing to locate performance improvements.

Our approach for using mutation to highlight performance improvements is general in that we do not target any specific type of code nor recommend any type of solution. Our approach could be augmented to utilise many of the known improvement techniques as listed in other work [30]. Our approach is designed for automated search and therefore may not be directly beneficial for use by human programmers. Our approach is instead expected to work well with automated program improvement approaches where mutation and testing are performed. We speculate that it may also be more generally applicable for different types of "performance". The basic idea is that by modifying code you can measure changes in the characteristic of interest. There is the potential to apply this approach to find code which has a big impact on memory, network or disk usage provided these characteristics can be measured.

REFERENCES

[1] A. Nistor, T. Jiang, and L. Tan, "Discovering, reporting, and fixing performance bugs," in *Proceedings of the 10th Working Conference on Mining Software Repositories*. IEEE Press, 2013, pp. 237–246.

[2] J. A. Jones and M. J. Harrold, "Empirical evaluation of the tarantula automatic fault-localization technique," in *IEEE/ACM International Conference on Automated Software Engineering*. ACM, 2005, pp. 273–282.

[3] E. Coppa, C. Demetrescu, and I. Finocchi, "Input-sensitive profiling," in *ACM SIGPLAN Notices*, vol. 47, no. 6. ACM, 2012, pp. 89–98.

[4] D. E. Knuth, "Structured programming with go to statements," *ACM Computing Surveys (CSUR)*, vol. 6, no. 4, pp. 261–301, 1974.

[5] M. Selakovic and M. Pradel, "Performance issues and optimizations in javascript: an empirical study," in *Proceedings of the 38th International Conference on Software Engineering*. ACM, 2016, pp. 61–72.

[6] R. Mudduluru and M. K. Ramanathan, "Efficient flow profiling for detecting performance bugs," in *Proceedings of the 25th International Symposium on Software Testing and Analysis*. ACM, 2016, pp. 413–424.

[7] D. Shen, Q. Luo, D. Poshyvanyk, and M. Grechanik, "Automating performance bottleneck detection using search-based application profiling," in *ISSTA*, M. Young and T. Xie, Eds. ACM, 2015, pp. 270–281. [Online]. Available: http://dl.acm.org/citation.cfm?id=2771783

[8] G. Jin, L. Song, X. Shi, J. Scherpelz, and S. Lu, "Understanding and detecting real-world performance bugs," in *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'12)*. Beijing, China: ACM Press, Jun. 2012, pp. 77–88.

[9] W. Weimer, T. Nguyen, C. Le Goues, and S. Forrest, "Automatically finding patches using genetic programming," in *International Conference on Software Engineering (ICSE)*. IEEE Computer Society, 2009, pp. 364–374.

[10] W. B. Langdon and M. Harman, "Optimising existing software with genetic programming," *IEEE Transactions on Evolutionary Computation*, 2013.

[11] F. Wu, W. Weimer, M. Harman, Y. Jia, and J. Krinke, "Deep parameter optimisation," in *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*. ACM, 2015, pp. 1375–1382.

[12] W. B. Langdon and J. Petke, "Software is not fragile," in *Complex Systems Digital Campus E-conference, CS-DC'15*, ser. Proceedings in Complexity, P. Bourgine and P. Collet, Eds. Springer, Sep. 30-Oct. 1 2015, p. Paper ID: 356, invited talk, Forthcoming. [Online]. Available: http://cs-dc-15.org/

[13] J. Gosling, B. Joy, G. Steele, G. Bracha, and A. Buckley, "The java language specification," https://docs.oracle.com/javase/specs/jls/se7/html/index.html, Feb. 2013.

[14] The Eclipse Foundation, "Java development tools," http://www.eclipse.org/jdt/, Nov. 2012.

[15] M. Kuperberg, M. Krogmann, and R. Reussner, "ByCounter: Portable Runtime Counting of Bytecode Instructions and Method Invocations," in *Workshop on Bytecode Semantics, Verification, Analysis and Transformation (European Joint Conferences on Theory and Practice of Software)*, 2008.

[16] B. Cody-Kenny, E. G. Lopez, and S. Barrett, "locoGP: improving performance by genetic programming java source code," in *Genetic Improvement 2015 Workshop*, W. B. Langdon, J. Petke, and D. R. White, Eds. Madrid: ACM, 11-15 Jul. 2015, pp. 811–818. [Online]. Available: https://www.scss.tcd.ie/~codykenb/locoGP.html

[17] B. Cody-Kenny and S. Barrett, "Performance localisation," *CoRR*, vol. abs/1603.01489v1, 2016. [Online]. Available: https://arxiv.org/abs/1603.01489v1

[18] Y. Jia and M. Harman, "An analysis and survey of the development of mutation testing," *IEEE transactions on software engineering*, vol. 37, no. 5, pp. 649–678, 2011.

[19] J. Petke, M. Harman, W. B. Langdon, and W. Weimer, "Using genetic improvement & code transplants to specialise a c++ program to a problem class," in *European Conference on Genetic Programming (EuroGP)*, 2014.

[20] K. Nagaraj, C. Killian, and J. Neville, "Structured comparative analysis of systems logs to diagnose performance problems," in *Presented as part of the 9th USENIX Symposium on Networked Systems Design and Implementation (NSDI 12)*, 2012, pp. 353–366.

[21] E. Schulte, Z. P. Fry, E. Fast, W. Weimer, and S. Forrest, "Software mutational robustness," *Genetic Programming and Evolvable Machines*, vol. 15, no. 3, pp. 281–312, 2012.

[22] M. J. Harrold, G. Rothermel, R. Wu, and L. Yi, "An empirical investigation of program spectra," *ACM SIGPLAN Notices*, vol. 33, no. 7, pp. 83–90, 1998.

[23] T. Ball, "The concept of dynamic analysis," in *Software Engineering—ESEC/FSE'99*. Springer, 1999, pp. 216–234.

[24] M. L. Vásquez, C. Vendome, Q. Luo, and D. Poshyvanyk, "How developers detect and fix performance bottlenecks in android apps," in *ICSME*, R. Koschke, J. Krinke, and M. P. Robillard, Eds. IEEE, 2015, pp. 352–361. [Online]. Available: http://ieeexplore.ieee.org/xpl/mostRecentIssue.jsp?punumber=7321954

[25] I. L. M. Gutiérrez, L. L. Pollock, and J. Clause, "SEEDS: a software engineer's energy-optimization decision support framework," in *ICSE*, P. Jalote, L. C. Briand, and A. van der Hoek, Eds. ACM, 2014, pp. 503–514. [Online]. Available: http://dl.acm.org/citation.cfm?id=2568225

[26] M. Grechanik, C. Fu, and Q. Xie, "Automatically finding performance problems with feedback-directed learning software testing," in *2012 34th International Conference on Software Engineering (ICSE)*. IEEE, 2012, pp. 156–166.

[27] C. Killian, K. Nagaraj, S. Pervez, R. Braud, J. W. Anderson, and R. Jhala, "Finding latent performance bugs in systems implementations," in *Proceedings of the eighteenth ACM SIGSOFT international symposium on Foundations of software engineering*. ACM, 2010, pp. 17–26.

[28] R. Impagliazzo, R. Paturi, and F. Zane, "Which problems have strongly exponential complexity?" in *Foundations of Computer Science, 1998. Proceedings. 39th Annual Symposium on*. IEEE, 1998, pp. 653–662.

[29] O. Olivo, I. Dillig, and C. Lin, "Static detection of asymptotic performance bugs in collection traversals," in *PLDI*, D. Grove and S. Blackburn, Eds. ACM, 2015, pp. 369–378. [Online]. Available: http://dl.acm.org/citation.cfm?id=2737924

[30] A. Nistor, P.-C. Chang, C. Radoi, and S. Lu, "C aramel: detecting and fixing performance problems that have non-intrusive fixes," in *Proceedings of the 37th International Conference on Software Engineering-Volume 1*. IEEE Press, 2015, pp. 902–912.

[31] A. Nistor, "Understanding, detecting, and repairing performance bugs," Ph.D. dissertation, University of Illinois at Urbana-Champaign, 2014.